Compositional Sentence Representation from Character within Large Context Text

Geonmin Kim, Hwaran Lee, Jisu Choi, Soo-young Lee
Korea Advanced Institute of Science and Technology, Daejeon, South Korea
{gmkim90, hwaran.lee, jisu916, sy-lee}@kaist.ac.kr

arXiv:1605.00482v3 [cs.CL] 3 Jun 2016

Abstract

This paper describes a Hierarchical Composition Recurrent Network (HCRN) consisting of a 3-level hierarchy of compositional models: character, word and sentence. This model is designed to overcome two problems of representing a sentence on the basis of a constituent word sequence. The first is a data sparsity problem when estimating rare words, and the other is no usage of inter-sentence dependency. In the HCRN, word representations are built from characters, thus resolving the data-sparsity problem, and inter-sentence dependency is embedded into sentence representation at the level of sentence composition. We propose a hierarchy-wise language learning scheme in order to alleviate optimization difficulties when training deep hierarchical recurrent networks in end-to-end fashion. The HCRN was quantitatively and qualitatively evaluated on a dialogue act classification task. Especially, sentence representations with inter-sentence dependency are able to capture both implicit and explicit semantics of sentences, significantly improving performance. In the end, the HCRN achieved state-of-the-art performance with a test error rate of 22.7% for dialogue act classification on the SWBD-DAMSL database.

1 Introduction

Sentence representations are usually built from representations of constituent word sequences by a compositional word model. Many compositional word models based on neural networks have been proposed, and have been used for sentence classification (Socher et al., 2010; Kim, 2014) or generation (Sutskever et al., 2014) tasks. However, learning to represent a sentence on the basis of constituent word sequences faces two difficulties.First, vector representation of word, word embedding, vary independently from other words. Thus, estimation of rare words suffer from the data-sparsity problem in which there are insufficient data samples to learn embedding of them. Poorly estimated of rare words embedding can cause sentence representations of inferior quality. Second, conventional sentence representation does not take into account the dependency of the meaning of one sentence on the meanings of other sentences. This inter-sentence dependency is especially evident in large context text such as in documents and in dialogues. Without accounting for this inter-sentence dependency, model capture only superficial meanings of sentences, and cannot capture implicit aspects of sentences such as intention which often requires linguistic context to understand.

In this paper, we propose a Hierarchical Composition Recurrent Network (HCRN), which consists of a hierarchy of 3 levels of compositional models: character, word and sentence. Sequences at each level are composed by a Recurrent Neural Network (RNN) which has shown good performance on various sequence modeling tasks. In the HCRN, output of lower levels of the compositional model is fed into the higher levels. Sentence representation by the HCRN enjoys several advantages when compared to sentence representation by a single compositional word model. From the compositional character model, the word representation

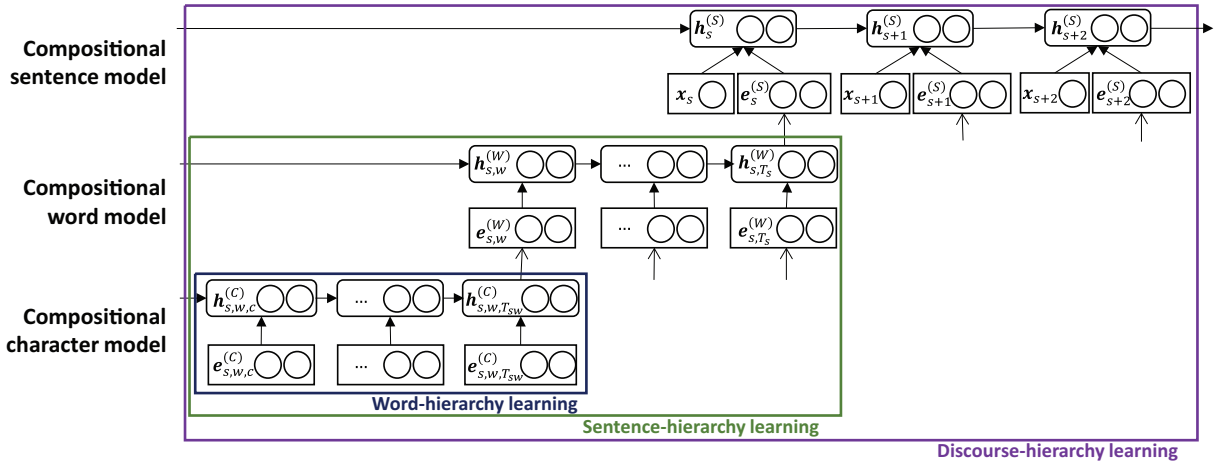

Figure 1: Illustration of the Hierarchical Composition Recurrent Network. The thick arrows indicate transformations described in equations (3)-(6), the thin arrows lines indicate identity transformation. For simplicity, each level is shown with one layer.

is built from characters by modeling morphological processes shared by different words. In this way, the data-sparsity problem with rare words is resolved. From the compositional sentence model, inter-sentence dependency can be embedded into sentence representation. Sentence representation with inter-sentence dependency is able to capture implicit intention as well as explicit semantics of given sentence.

Training the HCRN in an end-to-end has presents optimization difficulties since the network, a deep hierarchical recurrent network, may suffer from the vanishing gradient problem across different levels in the hierarchy. To alleviate this, hierarchy-wise language learning algorithm is proposed, and empirically shows that it improves the optimization of network. The hierarchy-wise learning algorithm trains the lower level network first, and then trains higher levels successively.

The efficacy of the proposed method is verified on a spoken dialogue act classification task. The task is to classify communicative intentions of sentences in spoken dialogues. Compared to conventional sentence classification, this task presents two challenging problems. First, it requires that the model estimate representations of spoken words which often include rare and partial words. Second, understanding dialogue context is often required to clarify meanings of sentences with a given dialogue.

The HCRN with hierarchy-wise learning algorithm achieves state-of-the-art performance on the SWBD-DAMSL database.

2 Hierarchical Composition Recurrent Network

Figure 1 shows our proposed Hierarchical Composition Recurrent Network (HCRN). The HCRN consists of a hierarchy of RNNs with compositional character, compositional word, and compositional sentence levels. At each level, each sequence is encoded by the output state of hidden layer of RNN at the end of sequence, which is commonly used in other research (Cho et al., 2014; Tang et al., 2015). In Figure 1 and in this section, each level of RNN is assumed to have one layer for simplicity of notation and figure. The notation of well-known transformations are represented as follows: gating units such as LSTM or GRU are represented as g, and affine transformation and non-linearity as r.

Consider a dialogue D which consists of sentence sequences $s_{1:T_D}$ and its associated label $t_{1:T_D}$. The compositional character model sequentially takes the c-th character embedding $\mathbf{e}^{(C)}_{s,w,c}$ of word w in sentence s, and recurrently calculates the output state of hidden layer $\mathbf{h}^{(C)}_{s,w,c}$ to produce a word representation $\mathbf{e}^{(W)}_{s,w}$ as:

$$\mathbf{h}^{(C)}_{s,w,c} = g(\mathbf{h}^{(C)}_{s,w,c-1}, \mathbf{e}^{(C)}_{s,w,c}) \quad (1)$$

$$\mathbf{e}^{(W)}_{s,w} = \mathbf{h}^{(C)}_{s,w,T_{sw}} \quad (2)$$

where T_{sw} is the length of character sequence of word w in sentence s.

Similarly, the compositional word model takes word sequence as input, and iteratively calculates the output state of hidden layer $\mathbf{h}_{s,w}^{(W)}$ to produce a representation of the s-th sentence $\mathbf{e}_s^{(S)}$ as:

$$\mathbf{h}_{s,w}^{(W)} = g(\mathbf{h}_{s,w-1}^{(W)}, \mathbf{e}_{s,w}^{(W)}) \tag{3}$$

$$\mathbf{e}_s^{(S)} = \mathbf{h}_{s,T_s}^{(W)} \tag{4}$$

where T_s is the number of words in sentence s.

The compositional sentence model updates its hidden neuron activation $\mathbf{h}_s^{(S)}$ in the same way as lower levels. Especially for dialogue, the compositional sentence model additionally takes the agent (or speaker) identity change vector $\mathbf{x}_s$, as an input. Agent identity, or at least agent identity change across neighboring sentences, is an important clue to understand the intended meaning of a sentence in a dialogue, as shown in previous research (Li et al., 2016).

$$\mathbf{h}_s^{(S)} = g(\mathbf{h}_{s-1}^{(S)}, \mathbf{e}_s^{(S)}, \mathbf{x}_s), \tag{5}$$

where agent identity change vector $\mathbf{x}_s$ is:

$$\mathbf{x}_s = \begin{cases} (1,0)^T, & \text{if } a_s = a_{s-1} \\ (0,1)^T, & \text{if } a_s \neq a_{s-1} \end{cases} \tag{6}$$

where a_s is the agent identity of s-th sentence.

The memory of the compositional sentence model $\mathbf{h}_s^{(S)}$, which includes s-th sentence information as well as that of previous sentences in the dialogue, is fed into a multi-layer Perceptron for classification of its label.

One advantage of the HCRN is its ability to learn long character sequences. While conventional stacked RNNs have difficulty when dealing with very long sequences (Bengio et al., 1994; Hochreiter, 1998), the hierarchy of the HCRN deals with short sequences of types specific to each level so that vanishing gradient problems during back-propagation through time are relatively insignificant. Each level of the HCRN uses a different speed of dynamics during sequence processing, so that the model can learn both short-range and long-range dependencies in large text samples. (See the preliminary experiment in supplementary materials of section A for details).

The following abbreviations are used for the rest of this paper: the compositional character model (CC), the compositional word model (CW), the compositional sentence model (CS), and the multi layer perceptron (MLP).

3 Hierarchy-wise Language Learning

In order to alleviate optimization difficulties occurring when entire hierarchies of RNNs are trained in an end-to-end fashion, hierarchy-wise language learning is proposed. In the hierarchy of composition models, the lower level composition network is trained first, gradually adding higher level composition networks after the lower level network is optimized for a given objective function. This approach is inspired by the unsupervised layer-wise pre-training algorithm in (Hinton et al., 2006), known to provide better initialization for subsequent supervised learning.

3.1 Unsupervised Word-hierarchy learning

To pre-train the CC, we adopt a pre-training scheme by (Srivastava, 2015), while following the architecture of RNN Encoder-Decoder in (Cho et al., 2014). Figure 2 shows the RNN Encoder-Decoder architecture used for this learning, and parameters of the CC can be obtained from the RNN Encoder.

In this architecture, the representation of word $\mathbf{e}^{(W)}$ which consists of characters $c_{1:T_W}$ is built by the CC. This representation is then fed into the RNN decoder. The CC and RNN Decoder are jointly trained so that their output sequence $\tilde{c}_{1:T_W}$ becomes exactly the same as the input character sequence $c_{1:T_W}$ by minimizing the negative log likelihood. This phase helps the CC learns *how to spell* words as character sequences, reducing the burden of learning morphological process of words on subsequent learning in higher levels of the hierarchy.

3.2 Supervised Sentence-hierarchy and Discourse-hierarchy learning

Next, sentence-hierarchy learning and discourse-hierarchy learning proceed in a supervised way with given sentence labels. In sentence-hierarchy learning, the model is trained to classify the label of single sentences independently from other sentences in a context. In discourse-hierarchy learning, the

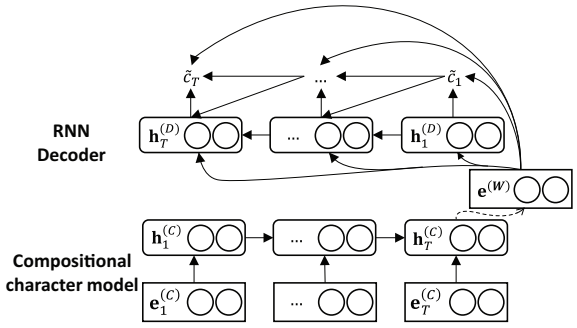

Figure 2: Illustration of the architecture used for word-hierarchy learning. Training objective forces the network to reconstruct exactly the same input character sequences.

model is trained to classify the label across multiple sentences. An MLP is stacked on top of sentence representation $\mathbf{e}_s^{(S)}$ (for sentence-hierarchy learning), or $\mathbf{h}_s^{(S)}$ (for discourse-hierarchy learning) to predict class labels.

A randomly initialized CW is stacked on top of the CC pre-trained during word-hierarchy learning, and a randomly initialized CS is stacked on top of the pre-trained CC and CW resulting from sentence-hierarchy learning. The parameters of the pre-trained lower level models are excluded from the training during the first few epochs in order to prevent the lower level pre-trained models from changing too much due to large error signals generated by the randomly initialized layer. This method is also employed in (Oquab et al., 2014) when adding a new layer on top of pre-trained layers for transfer learning. In section 4.5, the hierarchy-wise learning will be empirically shown to alleviates optimization difficulties in training the HCRN compared to end-to-end learning.

4 Experiments

4.1 Task and Dataset

The HCRN was tested on a spoken dialogue act classification task. The dialogue act (DA) is the communicative intention of a speaker for each sentence. Prediction of the DA can be further used as an input to modules in dialogue systems such as dialogue manager. We chose the SWBD-DAMSL database[1], which is a subset of the Switchboard-I (LDC97S62) corpus annotated with DA for each sentence. The Switchboard-I consists of phone conversations between strangers. The SWBD-DAMSL has 1155 dialogues of 70 pre-defined topics, 0.22M sentences, and 1.4M word tokens. There are several dialogue act tagsets available depending on the purpose of application. We chose the 42-class tagset from DAMSL, which is widely used to analyze dialogue acts of phone conversations. (See supplementary material section A for the complete class list.)

The number of elements in the character dictionary is 31 including 26 letters, - (indicating a partial word), '(indicating possessive case), . (indicating abbreviation), <noise> (indicating non-verbal sound) and <unk> (indicating unknown words) for all other characters. We follow the train/test set division in (Stolcke et al., 2000): 1115/19 dialogues, respectively. Validation data includes 19 dialogues chosen from training data. After pre-processing of the corpus, the number of sentences in the train/test/validation sets are 197370, 4190 and 3315 respectively. (See supplementary material section C for pre-processing details.) Table 1 shows the sequence length statistics. Note that a sentence becomes a much longer sequence when represented by characters (37.92) compared to words (8.28).

Table 1: Summary of the average sequence length for each level. C, W, S, and D indicate character, word, sentence and dialogue levels of processing, respectively. #A/B means the average number of A per B.

Hierarchy	#C/W	#W/S	#S/D	#C/S
Mean	8.19	8.28	161.26	37.92
Stddev	2.52	8.11	67.32	39.72

Table 2: Size of compositional model at each level, represented by (# layers) $\times$ (# memory in each layer). Note that the complexity of the model increases as the level of hierarchy increases, following the assumption that the complexity of composition increases as the level of language increases.

	CC	CW	CS
Small	1×64	2×128	2×256
Large	2×128	3×256	3×512

4.2 Common settings

We employ the Gated Recurrent Unit (GRU) as a basic unit of the RNN because several studies

[1]The dataset is available at `https://web.stanford.edu/~jurafsky/swb1_dialogact_annot.tar.gz`

have shown that the GRU performs similar with an LSTM, while requiring that fewer parameters be tuned (Chung et al., 2014; Jozefowicz et al., 2015). The configuration of the HCRN is represented by the hierarchy of the compositional model and its size, $CC_{size} - CW_{size} - CS_{size}$, where the size is represented by the number of layers and the number of hidden units in each layer. We tested two different sizes of compositional model at each level, as shown in Table 2.

In all supervised learning, the classifier consists of MLP with 3-layers of affine transformation and non-linearity. The non-linearity of the first 2 layer is Rectified Linear Unit (ReLU), and that of last layer is softmax. The number of hidden neurons used in MLP is 128. A common hyperparameter setting is used in all experiments. All weights are initialized from a uniform distribution within [-0.1, 0.1] except pre-trained weights. We optimized all networks with adadelta (Zeiler, 2012) with decay rate (ρ) 0.9 and constant (ϵ) 10^{-6} which show much faster convergence than stochastic gradient descent with momentum 0.9, and employed a gradient clipping strategy adopted from (Pascanu et al., 2013) with clipping threshold 5. Early stopping based on validation loss was used to prevent overfitting.

4.3 Unsupervised Word-hierarchy learning

During word-hierarchy learning, CC is jointly trained with the RNN Decoder to reconstruct input character sequences. The number of all unique words in the training set is 19353. The *end of word* token is appended at every end of character sequence. The parameters are updated after processing a minibatch of 10 words. The chosen character embedding dimension is 15. Learning is terminated if validation loss fails to decrease by 0.1% for three consecutive epochs.

Pre-training performance itself is evaluated by sequence reconstruction ability. For reconstruction, the RNN Decoder generates character sequences from the encoding $\mathbf{e}^{(W)}$ generated by CC. Generation is performed based on greedy sampling at each time step. The performance is evaluated on two measures: Character Prediction Error Rate (CPER) and Word Reconstruction Fail Rate (WRFR). CPER measures the ratio of incorrectly predicted characters to the reconstructed sequence. WRFR is the ratio of words where complete reconstruction fails out of the total words in test set.

Table 3: Reconstruction performance of the RNN Encoder-Decoder on words in the vocabulary and out-of-vocabulary (OOV). The length column presents the mean and the standard deviation (in parentheses) of character length of words for which complete reconstruction failed.

Model	In Vocabulary			Out of Vocabulary		
	CPER (%)	WRFR (%)	Length	CPER (%)	WRFR (%)	Length
$CC_{1\times64}$	0.39	2.25	13.1(2.6)	2.06	9.17	12.3(2.2)
$CC_{2\times128}$	0	0	-	1.21	5.28	12.7(2.4)

The reconstruction performance of the RNN Encoder-Decoder on words both in vocabulary and out-of-vocabulary (OOV) is summarized in Table 3. Overall, the model almost perfectly reconstructs character sequences of training data, and even generalizes well for the unseen words. The large size model outperforms the small size model. Almost all cases in which reconstruction failed involved sequences longer than 12 characters on average.

4.4 Supervised Sentence-hierarchy learning

During sentence-hierarchy learning, the parameters are updated after processing a minibatch of 64 sentences.

Initialization of CC: Random VS. Pre-trained The test set classification error rate of sentence-hierarchy learning with and without the pre-trained CC are compared to evaluate how the pre-trained CC provides useful initialization for sentence-hierarchy learning. With the pre-trained CC, at first, the parameters of the CC are frozen, and the CW and MLP are trained for 1 epoch[2]. After that, the whole architecture consisting of the CC, CW, and MLP is jointly trained. Evaluation was performed on architectures with different CC and CW sizes (see Table 2). In addition, pre-training on two different training dataset sizes (50 % and 100 %) are compared. The results are shown in Figure 3. Pre-training consistently reduces the test error rate on the various architectures as shown in Fig. 3(a). Moreover, as shown in Fig. 3(b), improvement by pre-training is significant especially when fewer train-

[2]The number of epochs to freeze the pre-trained model is chosen as the best parameter from preliminary experiments on the validation set.

Table 4: Comparison of word representation built from CC and non-compositional word embedding. The nearest 3 words by Euclidean distance are retrieved for given target word.

Word representation method	Unigram counts $\geqq 5$		$1 \leqq$ Unigram counts $\leqq 4$		Out of Vocabulary (OOV)	
	uh-huh	really	emphasizing	probab-	environmentalism	seventy-eights
Compositional Character Model	uh-oh huh-uh um	reall- real very	emphasize emphasis surpassing	probably probability probable	environmentalist environmentals environmental	ninety-eight seventeenth twentiy-six
Non-compositional ($\tau_c = 5$)	hmm helpful yeah	believe very frankly	-		-	

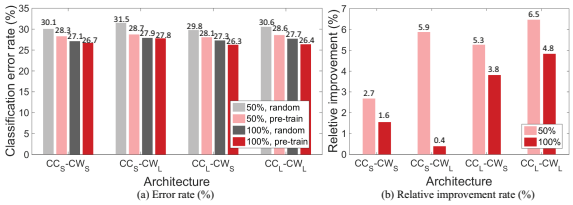

Figure 3: Result to show the quality of our pre-trained CC for initialization on sentence-hierarchy learning. (a) Test error rate (%) (b) Relative error improving rate of the pre-trained architecture with respect to the randomly initialized architecture. At each level, size is represented as either small (S) or large (L) in Table 2.

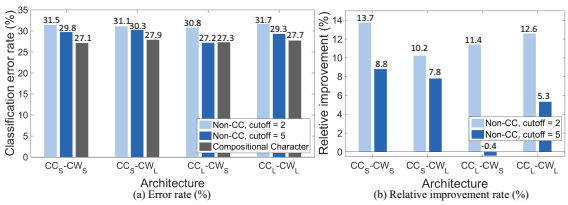

Figure 4: Result to show the quality of word representation built by CC. (a) Test error rate (%) (b) Relative error improvement rate of the CC with respect to non-compositional word representation with two different cutoff frequencies. At each level, size is represented as either small (S) or large (L) listed in Table 2.

ing data are available, where the model is liable to overfit to a given small number of data.

CC VS. non-compositional word embedding We compare two different methods to build word representation in this section: CC and non-compositional word embedding. Note that pre-trained CC are not directly compared to other widely used pre-trained word embedding such as Word2Vec (Mikolov et al., 2013) since pre-training of the CC aims at the model learning underlying morphological structures of words rather than learning the semantic/syntactic similarities between different words. Therefore, for a fair comparison we randomly initialized both models rather than employing pre-trained word embedding.

For the non-compositional word embedding method, we set two different cutoff frequencies: $\tau_c = 5$ (6294 words), $\tau_c = 2$ (11746 words). The dimensions of the non-compositional word embedding are chosen as 64 and 128, which is the same dimensions produced by sizes of CC in Table 2.

Figure 4 shows comparison of test error rates on above settings. Non-compositional word embedding with the high cutoff setting ($\tau_c = 5$) outperforms the low cutoff setting ($\tau_c = 2$). This is because the data sparsity problem during estimation of rare words is more severe for the model with the lower cutoff setting. Compared to the non-compositional method, CC outperforms or is on par, with fewer parameters since the number of parameters for non-compositional word embedding scales with the number of words in a dictionary.

Table 4 shows the 3-nearest neighbors of word representation built by two different methods: CC and non-compositional representation. For each method, the model with the best test accuracy is chosen. For word representation built by CC, retrieved nearest words usually have similar analogy with similar meaning. Moreover, rare words and OOVs such as partial words can be mapped to semantically similar words, which are usually estimated by a single <unk> in non-compositional word representation.

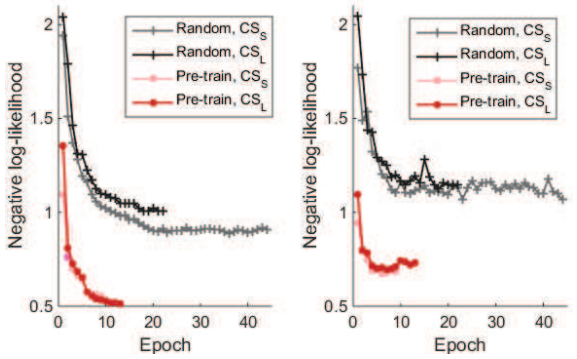

(a) Training loss (b) Test loss

Figure 5: Learning curve on (a) training data and (b) test data. The objective function is converged to much lower value when the model employ initialization from the pre-trained model resulting from sentence-hierarchy learning.

4.5 Discourse-hierarchy supervised learning

During discourse-hierarchy learning, the CS on top of the $CC_{2\times128} - CW_{2\times128}$ is trained. The pre-trained model was chosen by selecting the model with the lowest validation error rate during sentence-hierarchy learning. For the first 5 epochs, the network is trained with the CC and CW frozen. Then, the whole network is jointly optimized. The model is updated after it processes a minibatch of 8 dialogues. We evaluate performance of discourse-hierarchy learning with the two different sizes of CS listed in Table 2.

Optimization difficulty of end-to-end learning To verify that hierarchy-wise learning actually alleviates optimization difficulies, we compared the objective function curves of discourse-hierarchy learning using two different model initializations: with the pre-trained model from sentence-hierarchy learning, and with random initialization (end-to-end learning). The learning curve in Figure 5 clearly shows that initializing with the pre-trained model significantly alleviates optimization difficulties.

Effects of dialogue context on sentence representation Table 5 shows the test classification error rate of sentence-hierarchy learning and discourse-hierarchy learning. Compared to sentence-hierarchy learning, discourse-hierarchy learning improves performance significantly (up to 13.48% relative improvement of error).

To qualitatively analyze the improvement, we show examples of sentences on the test set for which prediction is improved by dialogue context. Analysis is done with model $CC_{2\times128} - CW_{2\times128} - CS_{2\times256}$, which achieved the best test accuracy during discourse-hierarchy learning. Table 6 shows a dialogue segment example including 8 sentences, with the predicted labels for each sentence from both sentence-hierarchy and discourse-hierarchy learning. Highlighted sentences indicate cases where discourse-hierarchy learning correctly predicts while sentence-hierarchy learning fails to predict. For example, "yeah" in the 3rd sentence of the example can be interpreted as both Agreement and Backchannel, and an informed decision between the two is only possible when the dialogue context is available. This example demonstrates that sentence representation with dialogue context helps to distinguish confusing dialogue acts (see supplementary materials section D for further examples).

Table 5: Test error rate of sentence hierarchy learning and discourse hierarchy learning. The relative improvement rate of error (Rel.) with respect to sentence-hierarchy learning are also provided.

Hierarchy	Model	Err (%)	Rel (%)
Sentence	$CC_{2\times128} - CW_{2\times128}$	26.27	-
Discourse	$CC_{2\times128} - CW_{2\times128} - CS_{2\times256}$	**22.73**	**13.48**
	$CC_{2\times128} - CW_{2\times128} - CS_{3\times512}$	22.99	12.49

Comparison with other methods Several other methods for dialogue act classification are compared with our approach in Table 7. Our approach outperforms the other benchmarks, achieving 22.7% classification error rate on the test set. Similar approaches employ a neural network based model which hierarchically composes sequences starting from word sequences (Kalchbrenner and Blunsom, 2013; Ji et al., 2016; Serban et al., 2016). We conjecture that the improvement demonstrated by our model is due to two factors. First, our model build word representations from constituent characters and so suffers less from the data sparsity problem when confronted with rare words, and second the hierarchy-wise language learning method alleviates optimization difficulties of the deep hierarchical recurrent network. To the best our knowledge, our model achieves the state-of-the-art performance of dialogue act classification on the SWBD-DAMSL database.

Table 6: An example of dialogue segment containing 8 sentences. Predictions of label from model of sentence-hierarchy learning (without dialogue-context) and discourse-hierarchy learning (with dialogue-context) are provided along with true labels.

Dialogue segment (with speaker)	True label	Estimated label without context	Estimated label with context
A: and uh quite honestly i just got so fed up with it i just could not stand it any more	Statement	Statement	Statement
B: is that right	**Backchannel-question**	**Yes-No question**	**Backchannel-question**
A: yeah	**Agreement**	**Backchannel**	**Agreement**
A: i mean this is the kind of thing you look at	**Statement**	**Opinion**	**Statement**
B: yeah	Backchannel	Backchannel	Backchannel
A: you sit there	Statement	Statement	Statement
A: and when you are writing up budgets you wonder okay how much money do we need	**Statement**	**Wh-question**	**Statement**

Table 7: Performance comparison with other methods for dialogue act classification on SWBD-DAMSL.

Method	Test err. (%)
Class based LM + HMM (Stolcke et al., 2000)	29.0
RCNN (Kalchbrenner and Blunsom, 2013)	26.1
HCRN with word as basic unit* (Serban et al., 2016)	24.9
Utterance feature + Tri-gram context + Active learning + SVM (Gambäck et al., 2011)	23.5
Discourse model + RNNLM (Ji et al., 2016)	23.0
Discourse-hierarchy learning	**22.7**

*This performance was evaluated by ourselves due to task difference.

5 Related works

The difficulty for RNNs learning long-range dependencies within character sequences has been addressed in (Bojanowski et al., 2016). Hierarchical RNNs have been proposed as one possible solution. (Graves, 2012; Chan et al., 2015) proposed sub-sampling sequences hierarchically to reduce sequence length at higher levels. (Koutnik et al., 2014; Chung et al., 2015) proposed a RNN architecture wherein different layers learn at different speeds of dynamics. Compared with these models, the HCRN deals with shorter sequences at each level, and thereby the vanishing gradient problem is rendered relatively insignificant.

There are several recent studies on representing large context text hierarchically for sentence (Li et al., 2015; Serban et al., 2016) and document classification (Tang et al., 2015). These approaches benefit from hierarchical representations which represents long sequences as a hierarchy of shorter sequences. However, the basic unit used in these approaches is the word, and models that begin at this level of representation open themselves to the data sparsity problem. This problem is somewhat resolved by building word representation from constituent sequences. Successful examples can be found in POS classification (Santos and Zadrozny, 2014) and language modeling (Botha and Blunsom, 2014; Ling et al., 2015; Kim et al., 2016).

6 Conclusion

In this paper, we introduced the Hierarchical Composition Recurrent Network (HCRN) model consisting of a 3-level hierarchy of compositional models: character, word and sentence. The inclusion of the compositional character model improves quality of word representation especially for rare and OOV words. Moreover, the embedding of inter-sentence dependency into sentence representation by the compositional sentence model significantly improves performance of dialogue act classification. This is because intentions which may remain ambiguous in single sentence samples are revealed in dialogue context, facilitating proper classification. The HCRN is trained in a hierarchy-wise language learning fashion, alleviating optimization difficulties with end-to-end training. In the end, the proposed HCRN using the hierarchy-wise learning algorithm achieves state-of-the-art performance with a test classification error rate of 22.7 % on the dialogue act classification task on the SWBD-DAMSL database.

Future work aims at the learning of hierarchy information of given sequential data without explicitly given hierarchy information. Another direction involves applying the HCRN to other tasks which

might benefit from OOV-free sentence representation in large contexts such as document summarization.

References

[Bengio et al.1994] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning Long-Term Dependencies with Gradient Descent is Dicfficult. *IEEE Transactions on Neural Networks*, 5(2):157–166.

[Bojanowski et al.2016] Piotr Bojanowski, Armand Joulin, and Tomas Mikolov. 2016. Alternative structures for character-level RNNs. In *arXiv preprint arXiv : 1511.06303*.

[Botha and Blunsom2014] Jan a. Botha and Phil Blunsom. 2014. Compositional Morphology for Word Representations and Language Modelling. In *Proceedings of the 31st International Conference on Machine Learning*.

[Chan et al.2015] William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. 2015. Listen, attend and spell. In *arXiv preprint arXiv : 1508.01211*.

[Cho et al.2014] Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734.

[Chung et al.2014] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. In *arXiv preprint arXiv : 1412.3555v1*.

[Chung et al.2015] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2015. Gated feedback recurrent neural networks. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, volume 37, pages 2067–2075.

[Gambäck et al.2011] Björn Gambäck, Fredrik Olsson, and Oscar Täckström. 2011. Active Learning for Dialogue Act Classification. In *Proceedings of Interspeech 2011*.

[Graves2012] Alex Graves. 2012. *Supervised Sequence Labelling with Recurrent Neural Networks*. Springer-Verlag Berlin Heidelberg.

[Hinton et al.2006] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554.

[Hochreiter1998] Sepp Hochreiter. 1998. The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6:107–116.

[Ji et al.2016] Yangfeng Ji, Gholamreza Haffari, and Jacob Eisenstein. 2016. A Latent Variable Recurrent Neural Network for Discourse Relation Language Models. In *arXiv preprint arXiv : 1603.01913*.

[Jozefowicz et al.2015] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. 2015. An Empirical Exploration of Recurrent Network Architectures. In *Proceedings of the 32nd International Converenfe on Machine Learning (ICML)*, pages 171–180.

[Kalchbrenner and Blunsom2013] Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent Convolutional Neural Networks for Discourse Compositionality. In *ACL WS on Continuous Vector Space Models and their Compositionality*, pages 119–126.

[Kim et al.2016] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. 2016. Character-Aware Neural Language Models. In *Proceedings of Association for the Advancement of Artificial Intelligence*.

[Kim2014] Y Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Nautral Language Processing (EMNLP)*, pages 1746–1751.

[Koutnik et al.2014] Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. 2014. A Clockwork RNN. In *Proceedings of The 31st International Conference on Machine Learning (ICML)*, volume 32, pages 1863–1871.

[Li et al.2015] Jiwei Li, Minh-Thang Luong, and Daniel Jurafsky. 2015. A Hierarchical Neural Autoencoder for Paragraphs and Documents. In *arXiv preprint arXiv : 1506.01057*.

[Li et al.2016] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A Persona-Based Neural Conversation Model. In *arXiv preprint arXiv : 1603.06155*.

[Ling et al.2015] Wang Ling, Tiago Luis, Luis Marujo, Ramon Fernandez Astudillo, Silvio Amir, Chris Dyer, Alan W Black, and Isabel Trancoso. 2015. Finding Function in Form: Compositional Character Models for Open Vocabulary Word Representation. In *Proceedings of Empirical Methods on Natural Language Processing (EMNLP)*, pages 1520–1530.

[Mikolov et al.2013] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems (NIPS)*.

[Oquab et al.2014] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks. In *Proceedings of the IEEE*

Conference on Computer Vision and Pattern Recognition.

[Pascanu et al.2013] Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2013. How to Construct Deep Recurrent Neural Networks. In *arXiv preprint arXiv : 1312.6026*.

[Santos and Zadrozny2014] CD Santos and B Zadrozny. 2014. Learning Character-level Representations for Part-of-Speech Tagging. *Proceedings of the 31st International Conference on Machine Learning*, ICML-14(2011):1818–1826.

[Serban et al.2016] Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. In *Special Track on Cognitive Systems at AAAI*.

[Socher et al.2010] Richard Socher, Christopher D Manning, and Andrew Y Ng. 2010. Learning Continuous Phrase Representations and Syntactic Parsing with Recursive Neural Networks. In *NIPS-2010 Deep Learning and Unsupervised Feature Learning Workshop*.

[Srivastava2015] Nitish Srivastava. 2015. Unsupervised Learning of Video Representations using LSTMs. In *Proceedings of International Conference of Machine Learning 2015*, volume 37.

[Stolcke et al.2000] A Stolcke, K Ries, N Coccaro, E Shriberg, R Bates, D Jurafsky, P Taylor, R Martin, C V Ess-Dykema, and M Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics*, 26(3):339–373.

[Sutskever et al.2014] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 3104–3112.

[Tang et al.2015] Duyu Tang, Bing Qin, and Ting Liu. 2015. Document Modeling with Gated Recurrent Neural Network for Sentiment Classification. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1422–1432.

[Zeiler2012] Matthew D Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method. In *arXiv preprint arXiv : 1212.5701*.

Supplementary material

A Effects of Hierarchical Composition

In order to learn character-level representation of large text such as sentence and dialogue, our model is hierarchically composed. As a preliminary experiment, we compared the performance of hierarchical (Fig 6-(a), sentence-hierarchy learning in paper) and non-hierarchical (Fig 6-(b), conventional stacked RNN) architectures on dialogue act classification tasks with the SWBD-DAMSL database.

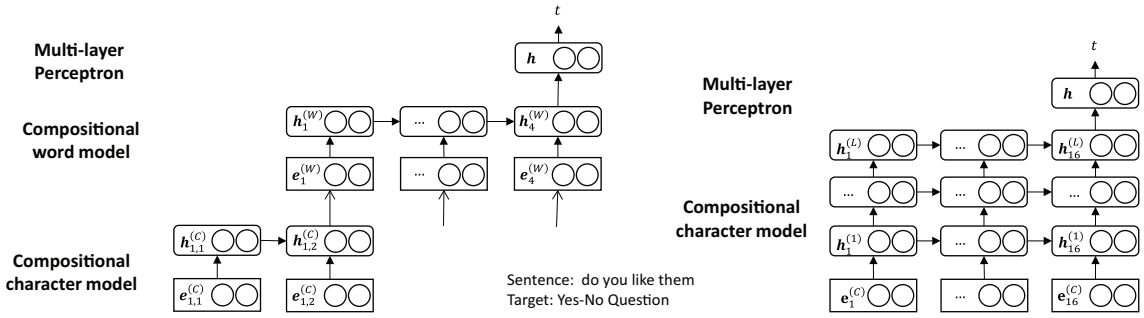

(a) Network with Hierarchical composition

(b) Network without Hierarchical composition

Figure 6: Comparison of network to learn sentence representation from character. (a) With hierarchical composition (b) Without hierarchical composition(conventional stacked RNN)

For fair comparison, hierarchical composition and non-hierarchical composition network have the same number of hidden units in each layer. For the non-hierarchical composition network, we add blank tokens in character sequences to indicate word boundaries. Table 8 shows that the hierarchical composition network outperforms the non-hierarchical composition network. It is because the network with hierarchical composition has a compositional word model (CW) which can be viewed as a specially designed compositional character model (CC), where composition is only allowed between hidden neurons at the ends of words. Therefore, the network with hierarchical composition has shorter sequences in each layer compared to the non-hierarchical model, where the vanishing gradient problem becomes relatively insignificant in each layer of RNN.

Table 8: Test error rate comparison of network with/without hierarchical composition. The relative improvement (Rel.) from non-hierarchical to hierarchical composition is also reported. + indicates consecutive layers in a conventional stacked RNN.

Non-hierarchical composition		Hierarchical composition		
Model	Error (%)	Model	Error (%)	Rel. (%)
$RNN_{1L,64H+2L,128H}$	35.80	$CC_{1L,64H} - CW2L, 128H$	27.13	24.22
$RNN_{1L,64H+3L,256H}$	36.38	$CC_{1L,64H} - CW3L, 256H$	27.81	23.56

B Complete list of class names in SWBD-DAMSL

Table 9: 42-class tagset of dialogue act provided from SWBD-DAMSL. Classes are sorted from the most frequent to the least frequent, from top-left to bottom-right with column-major order

Non-Opinion	Declarative question	Other answers
Backchannel	Backchannel(question)	Opening
Opinion	Quotation	Or clause
Abandoned	Summarize	Dispreferred answer
Agreement	Non-yes answer	3rd party talk
Appreciation	Action-directive	Offers
Yes-No-Question	Completion	Self talk
Non-verbal	Repeat phrase	Downplayer
Yes answer	Open question	Accept part
Closing	Rhetorical question	Tag question
Wh-question	Hold before answer	Declartive question
No answer	Reject	Apology
Acknowledgment	Non-no answer	Thanking
Hedge	Non-understand	Others

C Detail of pre-processing

We pre-processed raw text according to several rules below. At first, all letters are converted into lower-case. Disfluency tags and special punctuation marks such as (? ! ,) which cannot be produced by a speech recognizer are removed. We also merged sentences with segment tags into previous unfinished sentences. Segment tags indicate the interruption of one speaker by another. It is difficult even for human beings to predict tags of segmented sentences, because sentence segments often do not provide enough information for reliable ascription of its DA. Sentences which interrupt others are placed after combined sentences. This scheme is also used in (Webb et al., 2005; Milajevs and Purver, 2014).

D Supporting results showing discourse context improves sentence representation

D.1 Comparison of sentence-hierarchy learning and discourse-hierarchy learning on class accuracy

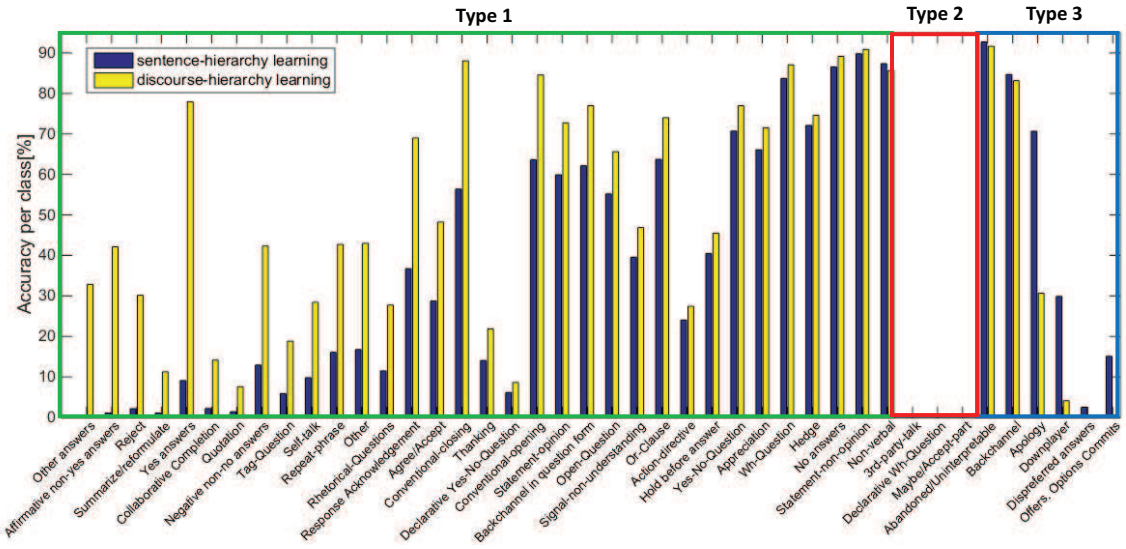

Figure 7: Class accuracy for sentence-hierarchy learning and discourse-hierarchy learning.

Figure 7 shows percent improvement when dialogue context is incorporated for each class label. Two models which achieves the best test set accuracy on sentence-hierarchy learning (without dialogue context) and dialogue-hierarchy learning (with dialogue context) are compared. Classes are sorted by descending order of relative improvement rate (from sentence-hierarchy to discourse-hierarchy). 33 out of 42 class improved with discourse-hierarchy learning (Type 1). Performance is degraded with discourse-hierarchy learning for 6 classes (Type 3). Also, there are 3 classes where both methods fail to predict. (Type 2)

D.2 Improved cases

Table 10: Examples of improved dialogue segment containing 8 sentences. More examples of Table 7.

Text	True label	Estimated label without context	Estimated label with context
B : what d- what is that	Wh-Question	Wh-Question	Wh-Question
A : it's more uh <noise>	Abandoned	Abandoned	Abandoned
A : i don't know how to explain it	Hold before answer	Opinion	Opinion
A : kind of pop you know rock	Statement	Statement	Statement
B : rock	**Repeat phrase**	**Statement**	**Repeat phrase**
A : yeah	**Agree**	**Backchannel**	**Agree**
B : hard rock	**Summarize**	**Statement**	**Summarize**
A : well not hard <laughter>	**Reject**	**Statement**	**Reject**
A : huh well are you going to paint the outside of your house too	**Yes-No-Question**	**Declarative question**	**Yes-No Question**
B : well yeah	Yes answers	Yes answers	Yes answers
B : i think i am going to do it this spring actually	Statement	Statement	Statement
A : oh really	Backchannel-question	Backchannel-question	Backchannel-question
B : yeah	**Yes answers**	**Backchannel**	**Yes answer**
B : there are six houses	Statement	Statement	Statement
B : see the people that own the house they uh pay for anything like that we do as far as the materials	**Statement**	**Opinion**	**Statement**
A : there are three houses on this street the same color of yellowout of six houses	Statement	Statement	Statement